



The State of Remote Scientific Visualization Providing Local Graphics Performance to Remote ARL MSRC Users

by John M. Vines and Claude Sandroff

ARL-TR-3635

September 2005

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5067

ARL-TR-3635**September 2005**

The State of Remote Scientific Visualization Providing Local Graphics Performance to Remote ARL MSRC Users

John M. Vines

Computational and Information Sciences Directorate, ARL

Claude Sandroff

IP Video Systems/Teraburst Networks Inc.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) September 2005		2. REPORT TYPE Final		3. DATES COVERED (From - To) March 2004–June 2005	
4. TITLE AND SUBTITLE The State of Remote Scientific Visualization Providing Local Graphics Performance to Remote ARL MSRC Users				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) John M. Vines and Claude Sandroff*				5d. PROJECT NUMBER 5U03CC	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: AMSRD-ARL-CI-HC Aberdeen Proving Ground, MD 21005-5067				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-3635	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES *IP Video Systems/Teraburst Networks Inc. Sunnyvale, CA					
14. ABSTRACT As data continue to grow beyond enormous sizes, visualization in three dimensions becomes the preferred way to analyze and solve complex problems. When visualization techniques are combined with technology that allows for global collaboration and sharing of high-resolution images, the Department of Defense has a tool for making faster and better decisions. Government and academic researchers engaged in analysis and simulation can access high-performance computing remotely from their desktops, improving both resource utilization and productivity. The need for networked visualization is strongly driven by those agencies and organizations where the cost or impact of decision making is very high. Such decision making revolves around large datasets, sophisticated modeling, and is based on complex data sharing and collaboration between remote heterogeneous teams. Clearly, organizations facing the need to analyze complex data sets will realize great benefits by creating networked visualization solutions.					
15. SUBJECT TERMS remote scientific visualization, IP communications, wavelet based compression, collaboration, visualization area network					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES 14	19a. NAME OF RESPONSIBLE PERSON John Vines
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED			19b. TELEPHONE NUMBER (Include area code) 410-278-9150

Contents

1. Introduction	1
2. Visualization and Collaboration	1
3. Networking Enhanced Visualization	1
4. Optimized Networking Solutions for Visualization	2
5. Visualized Images Contain a High Level of Redundancy	4
6. Compression Adapted for Visualization	4
7. Conclusion	6
Distribution List	7

INTENTIONALLY LEFT BLANK.

1. Introduction

As datasets continue to grow into the terabytes, users will be required to use distributed rendering and compositing techniques to enable the visualization of scientific data. Although users of the U.S. Army Research Laboratory Major Shared Resource Center will be able to take advantage of the latest distributed rendering and compositing techniques, remote users will still require a more viable connection avenue than XWindows can offer. The use of compression hardware and software will enable local graphics performance from a geographically dispersed location over an existing IP network.

2. Visualization and Collaboration

As data sets grow to enormous sizes, visualization in two or preferably three dimensions becomes the only way to analyze and solve complex problems. When visualization techniques are combined with technology that allows for collaboration and sharing of high-resolution images, Department of Defense agencies have a true competitive tool to foster collaboration and enhance productivity.

The ability to visualize and transport complex images in real time allows dispersed teams of experts in many geographically dispersed regions to make faster and better decisions.

Government and academic researchers engaged in analysis and simulation can access remote high-performance computing assets, improving both resource utilization and productivity.

As the “state-of-the-art” for data-visualization and data-assessment centers continues to evolve, visualization theaters and powerful computational facilities are continuously being built and upgraded. The need for networked visualization is strongly driven by those industries and organizations where the cost or impact of decision making is very high. Such decision making revolves around large datasets, sophisticated modeling, and is based on complex data sharing and collaboration between remote heterogeneous teams. Clearly, organizations facing the need to analyze complex data sets will realize great benefits by creating networked visualization solutions.

3. Networking Enhanced Visualization

Today’s high-end visualization typically uses SXGA screen resolution (1280×1024) but it is evolving towards UXGA (1600×1200), UXGA-W (1920×1200), and ultimately QXGA

(2048×1536). Display frame rates for viewing mono images most commonly range between 60 and 72 Hz, but range between 96 and 120 Hz for three-dimensional (3-D) stereoscopic imagery. The real rendered images, i.e., the different images generated by the visualization application, are 20–30 frames per s (fps), but it is expected that with high-end computers as well as clustered systems these rates will also go to 60 fps in the medium term. This means that the screen pixel content will increase from around 1 megapixel to more than 3.1 megapixels. At 24 bits of resolution per pixel (8 bits per color) and taking into account regular screen blanking, it means that uncompressed transmission rates will increase from 2 Gbps to close to 10 Gbps.

Adding networking to high-resolution visualization solutions allows organizations to address several major new application domains such as collaborative visualization, remote access, distribution, and broadcasting of high-end computer graphics to multiple locations simultaneously.

The main advantages of networked visualization are as follows:

- Sharing and distribution of high-resolution video data to multiple locations in the local, campus, and wide area. This allows organizations to keep the data set in a secure location while users access this data without having to travel to these highly secure data locations.
- Three-dimensional stereoscopic multiscreen graphics capability so that users can work with high-end visualization in any place, wherever the rendering computer is located.
- Security: the TeraBurst connectivity solutions only transport images in real-time, allowing the organization to keep all data in a secure centralized location. Hence, data at multiple sites no longer needs to be synchronized and updated, greatly simplifying data management. Furthermore, the public network never directly connects to the computer hosting the data, adding another layer of security. Images are transported only when a session is in progress, with a similar system being required at the remote end. Finally, circuit-based technologies like SONET/SDH tend to be highly secure while encryption systems, offered by IP Virtual Private Networks (VPN), can be added to packet-based solutions to provide full security.

4. Optimized Networking Solutions for Visualization

There are many solutions that allow networking of computers. Since we focus here on visualization requiring image transport with high-resolution and low latency, the following solutions are not considered:

- Solutions that are application, computer, and client dependant. Moreover, existing software solutions only address lower-end visualization needs (low resolution, two-dimensional, and/or low frame rate).

- Keyboard/video/mouse (K/V/M) hardware solutions (e.g., most of current packet-based solutions) that only address lower-end visualization, typically serves up to 1280×1024 at 60 Hz. They are mainly intended as remote consoles performing server management rather than visualization solutions.

The solutions discussed here (i.e., optical and high-end packet based solutions) are complimentary to these solutions serving medium- to high-end visualization needs.

TeraBurst has a deployed suite of digital optical products called Video-to-Optical (V₂O) that operate over both private fiber and public SONET/SDH networks. This digital approach—unique in the visualization market place—offers superior performance to current analog solutions for a number of reasons:

- They are standards-based and provide connectivity over virtually unlimited distances (tested up to 30,000 miles), safeguarding high (digital) quality with very low latency.
- They provide fully synchronized video, audio, and control signals (all are multiplexed over the same optical media) offering a fully interactive experience.
- They provide simple broadcast/multicast capabilities by splitting the signals to multiple destinations.
- They are easy to deploy since they only need one pair of single-mode fiber.
- They are secure since they do not transmit any “real data” and the multiplexed stream is complex to decipher. Moreover, optical monitoring tools can be used to further monitor any loss of optical signal intensity making the tapping of signals virtually impossible.

The TeraBurst packet-based (Ethernet/IP) solutions called Video-to-Data (V₂D) marry all of the major advantages of digital optical solutions with the ubiquity of data networks by the following:

- supporting short (LAN) and long (WAN) distances,
- providing a fully interactive experience through synchronized high-end video, audio, and control signals,
- offering broadcast/multicast capabilities through IP multicasting,
- using existing packet networks, and
- further enhancing security by adding data encryption (via IP VPN).

5. Visualized Images Contain a High Level of Redundancy

It is important to realize how much information is actually transmitted in most real applications. Uncompressed video translates into data rates that vary from 2 to 6 Gbps (and more). The second important parameter is the refresh rate vs. the real screen update rate. Typically, computers only render up to 20–30 fps while display rates can reach 96 Hz and higher. However, on the communication links, redundant frames can be eliminated by applying a number of techniques as follows:

- By systematically dropping frames (e.g., drop one or two frames every frame).
- By comparing whole frames and only sending new ones.
- By dropping identical “slices,” a slice defined as a well-defined part of a frame, between frames resulting in complete frame drop for identical frames and removal of additional redundant information in different frames

All these mechanisms dramatically reduce the required bandwidth with virtually no loss of information.

6. Compression Adapted for Visualization

The main goal of compression is to reduce the required bandwidth (over optical and packet networks) while safeguarding the overall picture quality so that the bandwidth cost becomes relatively insignificant when fielding visualization applications.

Off-the-shelf compression methods like MPEG2 and MPEG4 provide a cheap and highly compressed stream of video, but they do not fulfill the image quality requirements for high-end visualization. Therefore, specific hardware-based compression is required. This visualization compression makes use of reducing the redundant information and applying visually lossless compression in hardware to support high rendering and refresh rates.

Frame and slice differencing with variable noise reduction is the main mechanism to eliminate redundant information between frames. Noise reduction is a key mechanism, since many components introduce noise on the signals. Analog-to-digital conversion, in addition to system electronics, source stability (computer, video equipment), and intermediate (video) switches and cabling can all introduce noise. The level of noise is also largely dependent on the type of frames. For example, highly complex frames with sharp differences between pixels will be much more susceptible to noise than gradually changing images.

TeraBurst Networks has developed a compression solution (“TeraBurst Intelligent Frame Differencing^{*}”) which accounts for these requirements. When “static images” are displayed (e.g., analysis of a design) with an appropriate noise cancellation setting, close to zero information will be transmitted over the network. When “moving images” are displayed with the same noise cancellation setting, only new information will be transmitted.

Further reduction of bandwidth is achieved by applying spatial compression. As represented in figure 1, several measurements show that visually lossless compression can be implemented with a compression ratio up to 1:8–1:10 (depending on type of image, this requires 4-4-4[†] compression method). Higher compression rates up to 1:15 still provide good results, but some visual artifacts begin to be introduced. Also, when 4-2-2 compression is applied, visual artifacts become visible quicker but this compression method yields better results than 4-4-4 at higher compression ratios. Therefore, the recommended implementation method is a combination of both compression methods with user-selectability of the method and the corresponding parameters whereby the user can make a trade-off between bandwidth usage and image quality. The chart in figure 2 compares these two compression methods.

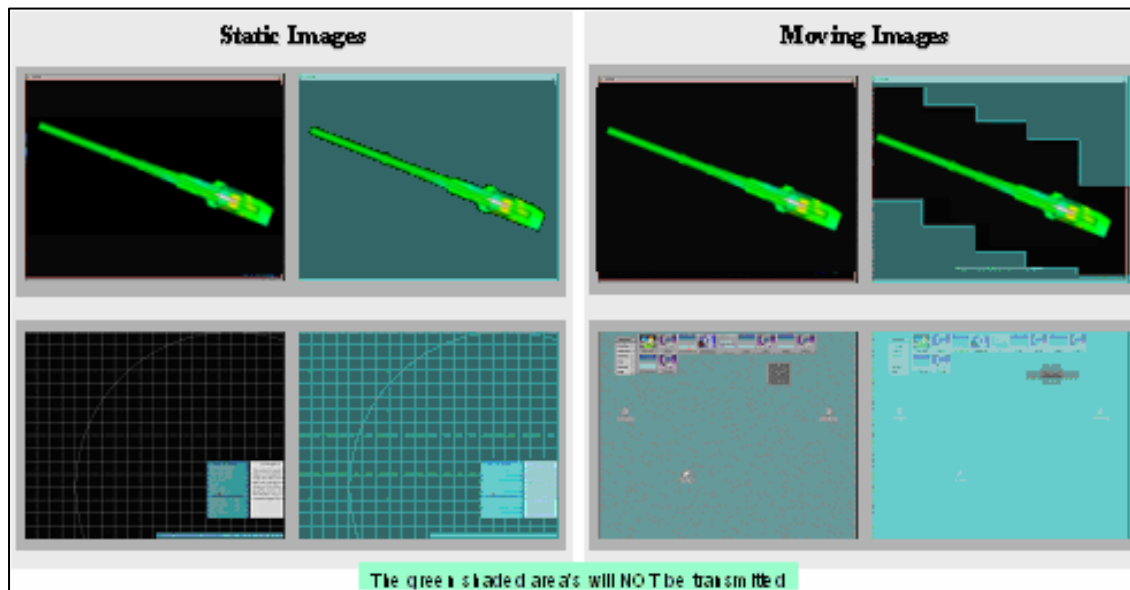


Figure 1. Static and moving images.

^{*}TeraBurst Intelligent Frame Differencing is a trademark of TeraBurst Networks, Sunnyvale, CA.

[†]4-4-4 compression uses as much luminance as chrominance information while 4-2-2 compression only uses half of the chrominance information.




Figure 2. Comparison of 4-4-4 and 4-2-2 compression methods.

With a combination of the compression methods, Gigabit (green) and Fast (yellow) Ethernet connectivity for high-end visualization can be implemented as shown in table 1.

Table 1. Transmission rate data.

Transmission Rate (Mbps)		Compression Slice Drop:	1:12 90%	1:8 75%	1:2 60%	Uncompression No Drop
Horizontal	Vertical	Refresh	(Mbps)	(Mbps)	(Mbps)	(Mbps)
1280	1024	60	16	59	378	1887
1280	1024	96	26	95	605	3020
1600	1200	60	24	87	554	2765
1600	1200	85	34	123	784	3917
2048	1536	85	55	202	1285	6417

7. Conclusion

The hardware compression equipment provided the remote user with real-time local graphics and file system performance where the traditional XWindows display provided an unusable interface to the data requiring the user to copy the data from a high-performance computing asset to the local file system to provide any means of interactivity.

NO. OF
COPIES ORGANIZATION

1 DEFENSE TECHNICAL
(PDF INFORMATION CTR
ONLY) DTIC OCA
8725 JOHN J KINGMAN RD
STE 0944
FORT BELVOIR VA 22060-6218

1 US ARMY RSRCH DEV &
ENGRG CMD
SYSTEMS OF SYSTEMS
INTEGRATION
AMSRD SS T
6000 6TH ST STE 100
FORT BELVOIR VA 22060-5608

1 INST FOR ADVNCD TCHNLGY
THE UNIV OF TEXAS
AT AUSTIN
3925 W BRAKER LN
AUSTIN TX 78759-5316

1 DIRECTOR
US ARMY RESEARCH LAB
IMNE ALC IMS
2800 POWDER MILL RD
ADELPHI MD 20783-1197

3 DIRECTOR
US ARMY RESEARCH LAB
AMSRD ARL CI OK TL
2800 POWDER MILL RD
ADELPHI MD 20783-1197

3 DIRECTOR
US ARMY RESEARCH LAB
AMSRD ARL CS IS T
2800 POWDER MILL RD
ADELPHI MD 20783-1197

ABERDEEN PROVING GROUND

1 DIR USARL
AMSRD ARL CI OK TP (BLDG 4600)

INTENTIONALLY LEFT BLANK.